

INCITE Proposal Preparation Instructions

2012 INCITE Call for Proposal

Average awards are expected to exceed 20 million processor hours. Requests for time below 10 million hours for science and engineering simulations may be redirected to another program.

Please note the new subsection 2d below. In CY2013 and beyond, the Leadership Computing Facilities anticipate putting into production next-generation resources. Proposals requesting multiple-year awards should address in this subsection their development plans to maximize node-level parallelism.

Section 4 is also new for the proposal template. All authors must submit a table that clearly articulates the scientific and technical milestones for each year of the proposed work.

Prior to submission, it is strongly recommended that proposers review their proposals to ensure that they comply with the proposal guidelines established below. These guidelines will be used to assist in the review of proposals. Templates for all sections are available. The INCITE program reserves the right to return without review any proposal that does not comply with the proposal guidelines.

The proposal must be clear, readily legible, and conform to the following requirements:

1. Each section of the proposal must be paginated. Footers should be used for paginating all files. Also, headers should be used to indicate title of the proposal and the lead PI.
2. Proposal titles may not exceed 80 characters in length.
3. Use one of the following typefaces identified below: Arial or Times New Roman (font size 12). A font size of less than 12 points may be used for mathematical formulas or equations, figure, table, or diagram captions and when using a symbol font to insert Greek letters or special characters. PIs are cautioned, however, that the text must still be readable.
4. Margins must be at least an inch in all directions. These requirements apply to all sections of a proposal, including supplementary documentation.
5. Proposals should be prepared using single line spacing.

Adherence to type size and spacing requirements also is necessary to ensure that no proposer will have an unfair advantage, by using smaller type or spacing to provide more text in the proposal.

The total file size should be limited to 5MB.

Proposal Contents

1. **Project Executive Summary (1 page):** This should include an executive summary that accurately describes the proposed research and the high-impact scientific or technical advances you will realize with the proposed INCITE allocation. Industry organizations should also summarize the potential economic or strategic business impact of the proposed research.
2. **Project Narrative: The narrative should not exceed 15 pages. Section limits are given below.** Visual materials, such as charts, graphs, pictures, etc will be included in the 15 page limitation. URLs that provide information related to the proposal should not be included in the proposal. The 15 page limitation will be strictly enforced. The Project Narrative should address the following points:
 - a. **Significance of Research:** Explain what advances you expect to be enabled by an INCITE award that justifies an allocation of petascale resources (e.g. anticipated impact on community paradigms, valuable insights into or solving a long-standing challenge, etc). Place the proposed research in the context of competing work in your discipline or business. List any previous INCITE award(s) received and discuss the relationship to the work proposed here. Please also list any previous Discretionary award(s) at the Argonne or Oak Ridge leadership computing facilities. **This section, including references, should not exceed four (4) pages.**
 - b. **Research Objectives and Milestones:** Describe the proposed research, including its goals, milestones and the theoretical and computational methods it employs. The information should be sufficient for peer review in your area of research and also appropriate for general scientific review comparing your proposal with proposals in other disciplines. Goals and milestones should articulate simulation and developmental objectives and be sufficiently detailed to assess the progress of the project for each year of any allocation granted. **This section, including references, should not exceed six (6) pages.**
 - c. **Computational Readiness: This section, including references, should not exceed five (5) pages.**
 - i. **Approach:** Provide
 1. a description of the underlying mathematical formulation (e.g., ODE, PDE),
 2. particular libraries required by your production and analysis software, the algorithms and numerical techniques employed (e.g., finite element, iterative solver), programming languages, and other software used,
 3. the parallel programming model(s) used (e.g., MPI, OpenMP, Pthreads), and

4. a description of project workflow including the role that analysis and visualization plays in the project's discovery process; identify where the analysis will be done and any potential bottlenecks in the analysis process.
- ii. **Job Characterization:** Applicants must present evidence that their proposed production simulations can make effective use of a significant fraction, in most cases 20% or more, of the HPC systems requested. Production jobs using <20% of the computer may experience significant throughput limitations. Contact the center to discuss their queuing policy if you anticipate difficulties. In your descriptions below state how the jobs are tied to each of your project's goals. For the simulations you plan to carry out during production runs, provide
 1. a description of what jobs are going to be run; relate these to the research/development objectives given above,
 2. a description of processor/core utilization for large runs (e.g. 10,000 hour run with 100 cores, or ten 10 hour runs with 10,000 cores, for a 1,000,000 hour allocation), and
 3. a clear, detailed explanation as to how you calculated the requested number of processor hours.
 - iii. **Parallel Performance:** Provide direct evidence, including *supporting quantitative data*, for your production application's parallel performance. Data for related – but not intended production – application codes is undesirable. Data should reflect performance for the intended research simulations. Performance benchmarking should reflect all I/O requirements of your application. Parallel performance data in either strong or weak scaling mode *must* be provided. Explain how the strong or weak scaling applies to the proposed work. Weak scaling behaviors are probed by holding per-node computational work constant (e.g., the size of the mesh on a processor is held constant) as the total problem size grows with processor count. Strong scaling behaviors are probed by holding the total problem size constant as the processor count grows, thereby decreasing the per-processor computational work. Supporting quantitative data should be provided in either tabular or graphical form, or both; also a speedup curve should be supplied for strong scaling examples. Where appropriate, characterize the production application's single-node performance (ex. percent of peak).

NOTE: If the supporting quantitative data is not available for your application, you may apply for a start-up account at one of the centers to conduct performance studies. Applications for start-up accounts are available at:

ANL: <http://www.alcf.anl.gov/support/gettingstarted/index.php>

ORNL: <http://www.nccs.gov/user-support/access/project-request/>

An example of the kind of parallel performance data sought can be found at the end of this document. Without this data, application readiness cannot be adequately ascertained and hence awarding this INCITE allocation request will be placed in jeopardy.

- iv. **Developmental Work:** If the performance data indicates less than 50% parallel efficiency on less than twenty percent of the requested system, please delineate the path forward for achieving greater performance on much larger scale simulations with the requested INCITE resources. If the production application scales (i.e. parallel speedup remains greater than 1) to less than 20% of the requested resource, or, if you propose to use part of an INCITE allocation to further enhance scaling and overall performance, outline the development plans and any desired support from the center. Describe what, if any, development work has been carried out to date, especially on the architecture of the requested resource.
- v. **Workflow Patterns:** Provide an estimate for the percentage of project time you will spend on development (porting, performance analysis) computing and other non-production runs. Describe your anticipated usage patterns for each year of the requested award, for example
 - Production jobs at a steady state throughout the year;
 - Development work at the beginning of the year, with periods [specified] of intense production work, or
 - Development work at the beginning of the year, with production work after mid-year.
- vi. **I/O:** Please discuss the I/O requirements for the proposed simulation jobs. Highlight any exceptional I/O needs.
 1. Restart I/O. Describe your I/O requirements for program restart. Include a list of the types of files you write out for program restarts, i.e. MPI/IO, raw binary, serial from all processors, HDF5, etc. and the following quantitative information: how much restart data you write out, or percentage of memory image occupied by the executing application, how many processors write-out restart data; how much memory your job takes per processor; the size of a single restart dump, maximum percentage of total execution time acceptable to read/write restart data, etc.
 2. Analysis I/O: describe your I/O requirements for analysis. Include a list of the types of files you write out for analysis (i.e. HDF5, NetCDF, PHDF5, PnedtCDF; size of analysis dump, amount and type of data read into analysis program, etc.), number of jobs needed on-hand to do analysis,
 3. Workflow I/O. How much I/O do you need per job? Do you have any exceptional I/O needs?
- vii. **Data storage:** The LCF reserves the right to request that data be transferred to another location following the project's completion. To

assist the INCITE program in identifying potential long-term storage solutions, please comments on the anticipated cumulative size of stored data at the end of the requested award (e.g. one-, two-, or three years) and any challenges in relocating this data.


- d. **Development plans for next-generation systems.** Proposals requesting time in CY2013 and beyond should provide a development plan articulating a strategy for running effectively on the anticipated next HPC systems at the ALCF and OLCF. See the system descriptions provided at the call for proposals web site (<http://hpc.science.doe.gov/>). We recognize that researchers may not have access to, or experience with, the new hardware architecture proposed by the LCF. **Assessment of development plans provided here will not affect decisions made for CY2012.** However, authors requesting time for CY2013 and beyond should give thought to the LCF future systems and indicate whether they have committed resources to prepare to maximize node-level parallelism. Areas to consider for discussion here include the following.
 - i. Hybridization plans utilizing OpenMP or Pthreads to expose thread-level or SMP-like parallelism
 - ii. Experiments or developments to expose vector or streaming parallelism through, e.g. CUDA, OpenCL, compiler directives, etc.
 - iii. General design to improve data locality and memory hierarchy usage
3. **Personnel Justification:** A personnel justification should be included in the proposal and detail what personnel are already in place, including a description of their roles on the project. If applicable, details should be provided on personnel that would be hired on the project in the future and the responsibilities of any new personnel. This should not exceed one page (Does not count towards the 15 page limit)
4. **Milestone Table:** New proposals must be accompanied by a summary table of planned milestones, listed for each year of the proposed work. Milestones should be clearly articulated and appropriate for the size and length of the requested award (e.g. large requests should have sufficient milestones to allow reviewers to assess the planned project workflow). Multi-year projects undergo an annual review and will be provided this table to update with project accomplishments. Any future modification to the project scope and milestones will be tracked in part through the milestone table. Reviewers of new proposals and renewals are specifically asked by the INCITE program about the milestones: the clarity and reasonableness of those in projects being considered for awards, and milestone achievements in projects being considered for renewal. It is important to spend sufficient time to clearly articulate the scientific and technical (e.g. development) milestones for each year of the proposed work.

NOTE: No Letters of Collaboration or Letters of Support will be accepted with the proposal. Current and Pending Support document is not required for this solicitation.

An example of the kind of parallel performance data sought is given in the following examples:

Weak Scaling Example:


| nProc | Time to Solution | Ideal |
|-------|------------------|-------|
| 4800 | 10.50 | 10.50 |
| 9600 | 10.45 | 10.50 |
| 19200 | 10.42 | 10.50 |
| 38400 | 10.40 | 10.50 |
| 76800 | 10.43 | 10.50 |



Strong Scaling Example:

| nProc | Time to Solution | Ideal |
|-------|------------------|---------|
| 4800 | 5000.00 | 5000.00 |
| 9600 | 2725.00 | 2500.00 |

| | | |
|-------|---------|---------|
| 19200 | 1500.00 | 1250.00 |
| 38400 | 1000.00 | 625.00 |
| 76800 | 700.00 | 312.5 |



Strong Scaling Speed Up Example:

| nProc | Speedup | Ideal |
|-------|----------|----------|
| 4800 | 4800.00 | 4800.00 |
| 9600 | 8807.34 | 9600.00 |
| 19200 | 16000.00 | 19200.00 |
| 38400 | 24000.00 | 38400.00 |
| 76800 | 34285.71 | 76800.00 |

Strong Scaling Speed Up Example

